An Oracle White Paper
December 2010

# NDMP Implementation Guide for the Sun ZFS Storage Appliance

ORACLE

Introduction ..... 1
Understanding Sun ZFS Storage Appliance platforms ..... 2
Network Data Management Protocol (NDMP) Overview ..... 3
Sun ZFS Storage Appliance NDMP Support ..... 3
Configuring NDMP on the Sun ZFS Storage Appliance ..... 4
NDMP Options ..... 4
Backup Considerations ..... 6
Attaching Tape Drives ..... 7
About ZFS-NDMP Backup Support ..... 8
Background ..... 8
Known Limitations ..... 8
Feature Overview ..... 8
Backup Application Configuration ..... 14
Oracle Secure Backup ..... 14
Symantec NetBackup ..... 17
EMC Legato Networker ..... 23
Conclusion ..... 28

## Introduction

With the large number of backup configurations and applications available today, it can be confusing how best to design a data protection architecture for important business data. This paper describes the implementation and configuration details for Network Data Management Protocol (NDMP) backups on the Sun ZFS Storage Appliance, using a variety of industry-leading backup applications:

- Oracle Secure Backup 10.3.0.2
- Symantec NetBackup 7.0
- EMC Legato Networker 7.6

In addition to application-specific configuration information this paper describes trade-offs typical in backup topologies and offers a new method of performing NDMP backups that is available with the Sun ZFS Storage Appliance 2010.Q3 software release.

## Understanding Sun ZFS Storage Appliance platforms

Sun ZFS Storage Appliances come in four platforms to meet the requirements for price/performance/capacity/protection capabilities. Sun ZFS Storage 7120 is entry level storage without read-optimized flash device. The other three platforms offer up to 2TB of read cache , which enables many applications to be on the read cache for substantially faster response time. The write flash on all the platforms is meant for providing quick response time to synchronous writes. With extremely fast CPUs of up to 8 threads per core, up to 32 threads are available to process the data in each appliance head. The mid-level appliance model Sun ZFS Storage 7320 offers 72GB of primary cache with mid-high range appliance models offer up to 512GB of memory. The new Sun ZFS Storage Appliances offer up to 2.5 TB of cache storage. From the application's perspective, this means more IOPS – especially for read-oriented applications.

**TABLE 1. SUN ZFS STORAGE PLATFORMS**

| PLATFORM | STORAGE CAPACITY | PROCESSOR | MEMORY (DRAM) | WRITE OPTIMZED SSD | READ OPTIMZED SSD | CLUSTER OPTIONS |
|---|---|---|---|---|---|---|
| Sun ZFS Storage 7120 | Up to 60 x 2TB SAS Disks [120TB] | 1 x Quad Core Intel Westmere EP E5620 @ 2.4GHz | Up to 36GB | Up to 98GB | N/A | N |
| Sun ZFS Storage 7320 [details are per controller] | Up to 96 x 2TB SAS Disks [192TB] | 2 x Quad Core Intel Westmere EP E5620 @ 2.4GHz | Up to 72GB | Up to 16 x 18GB | Up to 4 x 512GB | Y |
| Sun ZFS Storage 7420 [details are per controller] | Up to 576 x 2TB SAS Disks [1.1PB] | 4 x 6C Intel Nehalem EX E7530 @ 1.86GHz<br>Or<br>4 x 8C Intel Nehalem EX X7550 @ 2GHz | Up to 512GB | Up to 96 x 18GB | Up to 4 x 512GB | Y |
| Sun ZFS Storage 7720 | Expandable Racks. Each Rack 720 TB | 4 x 8C Intel Nehalem EX X7550 @ 2GHz | Up to 512GB per controller | 2 x 18GB per disk cage | Update to 4 x 512GB per controller | Y |

## Network Data Management Protocol (NDMP) Overview

NDMP is an open standard protocol for network-based backup of network-attached storage. It enables backup vendors and NAS storage vendors to work together without creating specific ports for each vendor and storage array in addition to allowing storage array vendors to have their arrays "backup-ready" without installing extra backup client software.

NDMP also allows for the direct attachment of tape or robotic libraries (including Virtual Tape Libraries) to the storage array, eliminating the need for backups to travel over a Ethernet network to a centralized backup server.

The Sun ZFS Storage Appliance supports NDMP versions 3 and 4. At the time of this writing, most backup applications have switched to NDMP version 4. The testing in support of this document was done using NDMP version 4.

## Sun ZFS Storage Appliance NDMP Support

The Sun ZFS Storage Appliance uses and recommends NDMP for backup and restore of the data sets, shares and block targets. It is possible to perform backups of Network File Systems (NFS) or Common Internet File Systems (CIFS) directly, without the use of NDMP, however this is not the preferred method.

NDMP offers the following advantages:

- Tape Drives can be attached, via a Fiber Channel SAN, directly to the Sun ZFS Storage Appliance for faster performance.
- All file system attributes, permissions, share settings and meta-data are backed up.
- The NDMP service on the appliance takes an automatic snapshot to ensure consistency of data.

The Sun ZFS Storage Appliance supports the following NDMP topologies:

- Local - The tape drive is connected to the appliance via Fiber Channel. Backup data moves directly to the tape drive without going over an Ethernet network to a backup server. This method usually provides the best performance. Depending on the backup application, the tape drive could be in a shared SAN environment and be made available to other appliances or even the backup server itself to perform other backups as needed.
- Remote - The tape drive or storage device (disk storage, Virtual Tape library) is connected to a backup media server. Data travels over a Ethernet Connection (1000BaseT or 10GB Ethernet recommended) to the backup server.
- 3-Way - A tape drive connected to one appliance is used to store data from another Sun ZFS Storage Appliance's backup. The data travels over an Ethernet network to the appliance controlling the tape resource. This is similar to a Remote backup.

Additionally, the following NDMP features are supported

- Direct Access Recovery
- Single File Restore (with dump or tar methods)
- ZFS-NDMP backup (See ZFS-NDMP Backup section for more details)
- Multi-Stream Concurrent Backups
- Incremental/Level Support

## Configuring NDMP on the Sun ZFS Storage Appliance

You configure the NDMP service for the Sun ZFS Storage Appliance from the Browser User Interface (BUI):

1. Log in to the Sun ZFS Storage Appliance BUI
2. Select the Configuration:Service:NDMP tab

### NDMP Options

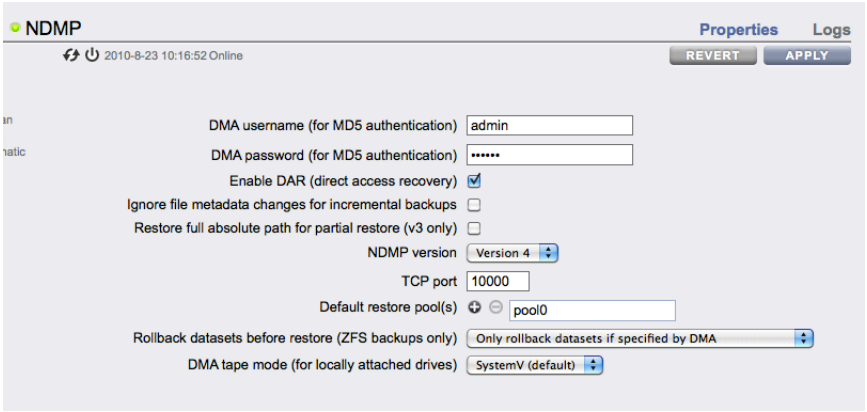


#### DMA Username and DMA Password

The administrator should provide a Data Management Application (DMA) username and DMA password. This username-password combination will also be used when configuring the backup application and must be consistent between the array and the application. The username and password do not need to be defined on the array as a Local User, nor do they need to be defined in a Directory Service such as NIS or Active Directory. This username and password are used only for NDMP operations.

#### Enable DAR

Enables the system to locate files by position rather than by sequential search during restore operations. Enabling this option reduces the time it takes to recover a small number of files from many

tapes. You must specify this option at backup time in order to be able to recover individual files later. In most situations, the default should be left on.

#### Ignore file metadata changes for incremental backups

Directs the system to back up only the files in which content has changed, ignoring files for which only metadata, such as permissions or ownership, has changed. This option only applies to incremental "tar" and "dump" backups and is disabled by default.

#### Restore full absolute path for partial restore (v3 only)

Specifies that when a file is restored, the complete absolute path to that file is also restored (instead of just the file itself). This option is disabled by default.

#### NDMP version

The version of NDMP that your DMA supports. The default is Version 4. Most backup applications support NDMP version 4.

#### TCP port

The NDMP default connection port is 10000. NDMP version 3 always uses this port. NDMP version 4 allows a different port if needed.

#### Default restore pool(s)

When doing a full restore using types "tar" or "dump," the system will re-create datasets if there is not already a share mounted at the target. Because the NDMP protocol specifies only the mount point, the system will, by default, choose a pool in which to recreate any projects and shares. For a system with multiple pools, specify one or more pools to restore to. Multiple pools need only be specified in a cluster with active pools on each head; it is the user's responsibility to make sure that this list is consistent with any storage configuration changes (such as adding new pool, or deleting a pool). If the pools are offline, or if none of the pools specified exist, then the system will select a default pool at random.

#### Rollback datasets before restore (ZFS backups only)

Only applies to zfs type backups. Determines whether, when restoring an incremental backup, the system rolls back the target project and share to the snapshot used as the base for the incremental restore. If the project and shares are rolled back, then any changes made since that snapshot will be lost. This setting is normally controlled by the backup application via the "ZFS_FORCE" environment variable (see the ZFS-NDMP section later in this paper) but this property can be used to override the backup application setting to *always* rollback these datasets or *never* roll them back. Not rolling them back will cause the restore to fail unless they have already been manually rolled back. This property is intended for use with backup applications that do not allow administrators to configure custom environment variables like ZFS_FORCE.

#### DMA tape mode (for locally-attached drives)

Specifies whether the backup application expects SystemV or BSD semantics. The default is SystemV, which is recommended for most backup applications. This option is only applicable for locally attached tape drives exported via NDMP. Consult your application documentation for which mode your application expects. Changing this option only changes which devices are exported when the backup application scans for devices, so you will need to reconfigure the tape devices in your backup application after changing this setting.

## Backup Considerations

The NDMP service on the Sun ZFS Storage Appliance accepts both dump and tar backup requests from the backup applications. These are the most common backup-type requests. The 2010.Q3.0 software for the Sun ZFS Storage Appliance adds a third backup type called "zfs" which leverages the ZFS file system capabilities to perform backups of block targets for Fiber Channel or iSCSI as well as offer performance improvements in backing up NFS or CIFS file systems. For more details on the ZFS-NDMP feature, see the "ZFS-NDMP Backup Support" Section below.

The rest of this section will focus on the default dump/tar methods of backup.

#### Backup Path

The NDMP service should only be used to back up user-created shares and file systems on the Sun ZFS Storage Appliance. This usually entails configuring the backup application to protect shares that are exported under "/export". Also, individual shares (or subdirs in shares) should be selected. Do not attempt to back up "/export" by itself.

**Correct backup paths**

/export/myshare

/export/myshare/subdir

**Incorrect backup paths**

ALL_LOCAL_DRIVES

Saveset = All

/export

/

/etc

/var

CAUTION: Trying to back up a root file system share will result in a failure and the backup will abort. To perform configuration backups of the array settings (network configuration, local users, roles, service settings), please see the "Configuration Backup/Restore" feature for the Sun ZFS Storage Appliance in the Online Help.

It is possible to create a Project on the Sun ZFS Storage Appliance that has a mount point of **/export/myproject**. All shares in that Project will inherit that base mount point and have a mount point of **/export/myproject/myshare1**, **/export/myproject/myshare2**. It is acceptable to create a backup that protects **/export/myproject** and all shares in that project that have the base **/export/myproject** mount point will be backed up in the same backup stream.

### Snapshots

Before a backup begins, the NDMP service on the Sun ZFS Storage Appliance takes a snapshot of the share being backed up to ensure a consistent, point-in-time image. After the backup is completed, this snapshot is deleted.

It is possible to take a backup of an existing snapshot that was created as part of a schedule or created as a manual, one-time snapshot. To backup an existing snapshot, use the ZFS snapshot path to the share.

#### Example: Snapshot backup

An administrator creates a snapshot named **end_of_year** for a share called **/export/accounting_data**. To back up this snapshot, the administrator configures the backup application to back up the path **/export/accounting_data/.zfs/snapshot/end_of_year**. Since the snapshot is already in a read-only, consistent state, the NDMP service does not create a separate backup snapshot and the existing snapshot is used as the data source for the backup.

## Attaching Tape Drives

The Sun ZFS Storage Appliance supports attaching Tape Libraries and Tape Drives via the optional Fiber Channel adapters. Normal Storage Area Network (SAN) configuration best practices should be followed to connect these devices and the fiber ports on the appliance:

- The devices should be configured as single-path only. Multiple paths to the devices are not supported.
- After attaching the devices, the appliance should be rebooted.

After the appliance boots, ensure the devices are seen by the NDMP service by viewing the Status:NDMP screen in the BUI.

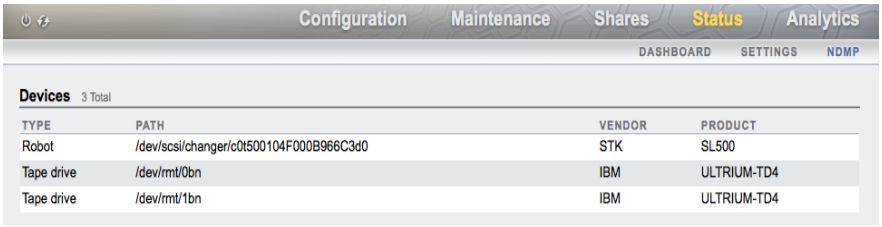

# About ZFS-NDMP Backup Support

The Sun ZFS Storage Appliance firmware release 2010.Q3 introduces a new NDMP backup methodology with the underlying method being ZFS send and receive. This new backup method can be used to back up block-based LUNs (such as iSCSI or Fiber Channel) and NFS/CIFS filesystems.

## Background

The Sun ZFS Storage Appliance NDMP service implements the Solaris NDMP service. This service provides backup and restore functionality per the NDMP protocol when used in conjunction with client DMAs. The NDMP service advertises support for both the "dump" and "tar" methods of backup and restore, even though both are internally implemented using "ustar." Additionally, the existing NDMP service cannot back up block-based LUNs such as Fiber Channel or iSCSI. The new ZFS-NDMP feature enables administrators to back up LUNs via NDMP, in addition to NFS/CIFS shares, as part of the overall data protection solution.

Though the existing NDMP service is able to back up NFS/CIFS shares using tar, special purpose code is required to back up underlying ZFS file system metadata. ZFS send and receive, on the other hand, is able to handle ZFS metadata transparently. In addition, it has the capability of backing up ZFS descendants of the dataset, including clones and snapshots. It can also preserve the sparseness of a file system and volume in the backup and restore process.

ZFS send and receive has been demonstrated to provide superior performance over directory traversal backup/restore. Such a result is due to the method's block traversal algorithm and its inherent knowledge of ZFS metadata layout and pre-fetching optimizations. It is expected that the use of ZFS send and receive will bring about superior performance in NDMP backup/restore environments without requiring significant change in infrastructure.

## Known Limitations

Current functionality of the ZFS-NDMP feature does not allow the cataloging and/or restore of individual files within any particular backup. At this time, it is not possible for a user to restore an individual directory or file that was backed up using this feature. All data contained the backed up dataset must be restored. Sufficient space on the Sun ZFS Storage Appliance must exist to restore the entire contents of the needed dataset prior to the restore being started.

**NOTE**: The ZFS-NDMP feature is only available with NDMP version 4.

## Feature Overview

### TYPE = ZFS

The ZFS send and receive functionality will be accessible in the NDMP service via a new backup type called "zfs." The ability to set this NDMP protocol variable, "TYPE=zfs," is backup application dependent. At the time of this paper's writing, the following backup applications had been tested and known to work with the TYPE=zfs setting.

- Oracle Secure Backup 10.3.0.2 and above
- Enterprise Backup Software (EBS) / Legato Networker 7.5 and above
- Symantec NetBackup 6.5.3 and above

Configuration examples for these three applications are given in **Backup Application Configuration** section of this paper.

Other backup applications, if they are capable of changing the backup type (TYPE=) variable at backup time, may also work.

#### ZFS_MODE

Two types of dataset backups are available under this "zfs" type. A dataset is a discrete NFS/CIFS share or block LUN (iSCSI or FC).

- dataset - Simple backup and restore of the indicated dataset.
- recursive - Recursive backup and restore, which includes all descendant datasets, but no intermediate snapshots.

These methods are specified using the ZFS_MODE environment variable. The default value is recursive. The ability to set this NDMP variable, "ZFS_MODE=[dataset | recursive ]," is backup application-dependent. See the "TYPE=zfs" section for backup applications that have been tested and known to work.

#### Other NDMP Environment Variable

There are other NDMP environment variables which you can use to control the behavior of the ZFS-NDMP backup feature:

- LEVEL - Indicates the incremental level for which the given snapshot will serve as endpoint. L0 indicates a full backup, and L1 - L9 indicate incremental backups, where only changes since the last level's backup are backed up. (default = 0)
- DMP_NAME - Allows a unique backup set name to be specified, such that a given dataset can be backed up in more than one incremental's set at a time. DMP_NAME will be validated to ensure that it includes alphanumeric characters only, with the exception of the underscore and hyphen characters. In addition, the length of DMP_NAME will be limited to 31 characters. (default = "level")
- UPDATE - Controls how snapshots are taken. (default = "y") See the "UPDATE Behavior" section for more details.
- ZFS_FORCE - Forces incremental restore (default = "n")

#### Backup Path

For "zfs" backup, the path consists of the ZFS dataset name. (Note that this is in contrast to traditional backup methods, that use a file system path to indicate the directory to be backed up.)

Paths specified in conjunction with the new "zfs" backup type will be parsed according to the following format (which is the same as ZFS dataset name syntax, except for the leading slash):

/<dataset>[@snap]

For example:

/pool_0/local/myproject/myfilesystem

/pool_0/local/myproject/mylun

The full dataset path is displayed in the Sun ZFS Storage Appliance BUI under the General Properties tab for each NFS/CIFS share or iSCSI/FC LUN.

The leading slash "/" is required.

As with dump/tar backup types, the ZFS backup method can also take backups of existing snapshots.

#### Example: Snapshot backup

An administrator creates a snapshot named **end_of_year** from a dataset with a ZFS backup path of **/pool_0/local/corporate/accounting**. To back up this snapshot the administrator would configure the backup application to back up the path **/pool_0/local/corporate/accounting@end_of_year**.

### Temporary Snapshots

If the backup path is not of a snapshot, a temporary snapshot will be taken and used for backup. If the NDMP environment variable UPDATE is set, this temporary snapshot will replace any existing temporary snapshot for the same LEVEL, DMP_NAME and ZFS_MODE configuration . The next time a temporary snapshot is taken for the same LEVEL, DMP_NAME and ZFS_MODE, this temporary snapshot will be replaced. If UPDATE is not set, the temporary snapshot will be deleted immediately after backup and any existing temporary snapshot for the same LEVEL, DMP_NAME and ZFS_MODE will be preserved. (Note that, in this case, it will not be possible for a subsequent incremental backup based on the current non-UPDATE backup to be performed.)

The syntax for the naming of this snapshot is as follows:

ndmp.0.level.r.1279023238.107462
ndmp.$LEVEL.$DMP_NAME.$ZFS_MODE.<timestamp>.

"ndmp" indicates that this snapshot was generated by the NDMP service. "0" indicates that this is a Level 0 (full) backup. "level" is the default DMP_NAME. "r" indicates a recursive backup as indicated by ZFS_MODE (other values could be "d" for dataset). The rest of the string is a system generated timestamp.

### Restore Path

The restore name should be of the form

/<dataset>[@<snap>]

For example:

/pool_0/local/myproject/myfilesystem-restore

As with the backup path, the restore dataset name should adhere to standard ZFS dataset name syntax.

When restoring from a full backup, the specified dataset name must not exist. To fulfill this requirement, the administrator might need to first remove or rename any dataset with a duplicate name. If the dataset does exist, the restore will fail and an error message will be logged to the backup application.

**NOTE**: All components except for the final component in the restore path must already exist prior to the restore. For example, if "/pool_0/local/myproject/myrestore" is specified as the nonexistent dataset to be used for restore, "/pool_0/local/myproject" must already exist prior to the restore operation, or an error will result.

If a snapshot name is specified at restore time, the restore process will ignore this extension and will perform the restore as if the snapshot were not specified. (Note that this is dependent on whether the backup application will allow the "@" character in the restore path. If not, the restore attempt may fail.) The advantage of this behavior is (depending on the backup application) it may allow the restore option, "Restore to original location," to work transparently (e.g.., if a snapshot had been specified at backup time). To mitigate any confusion, however, a message will be logged to backup application console and to the NDMP system log that the "@<snap>" extension is being ignored.

**NOTE**: "Restore to original location" is typically one of several options available to users, with "Restore to different location" also available.

### Incremental Backup

A backup of a ZFS file system using Level 0 (L0) will generate a full ZFS-NDMP backup stream. (This is the same as the "full" backup described elsewhere.) Levels 1-9 will generate incremental streams. For a given L incremental stream to be generated, an L-1 snapshot must already be present in the system. (This L-1 snapshot would have been generated by a prior L-1 backup.) Also, either the user must supply an L snapshot or the NDMP service will generate its own L snapshot. The NDMP service on

the Sun ZFS Storage Appliance will use the L-1 and L snapshots as endpoints for the incremental stream.

If a snapshot is indicated by the user for a given incremental level, that snapshot will be used for the ZFS-NDMP backup. If no snapshot is indicated on the command line, the NDMP service will take its own snapshot and name the snapshot according to the convention described in the "Temporary Snapshot" section of this paper.

**NOTE:** A snapshot created by NDMP should not be used subsequently as a backup input path. If it is used, the backup will fail. The reason for this requirement is to simplify the NDMP service's management of such snapshots. For instance, such snapshots are currently removed from the system based on the setting of the UPDATE variable.

### UPDATE Behavior

Each incremental backup (excluding Level 0 or full) requires two snapshots. One with value of $LEVEL-1.$DMP_NAME.$ZFS_MODE, and one with a value of $LEVEL.$DMP_NAME.$ZFS_MODE (i.e. a current snapshot). These snapshots are either user-supplied or ndmpd-generated.

When an incremental backup is done for a dataset, the value of the UPDATE environment variable determines whether any previously existing snapshot with the same $LEVEL.$DMP_NAME.$ZFS_MODE value remains accessible to the NDMP service for future backups. At the start of a backup, the snapshot to be used will have its property set with the value $LEVEL.$DMP_NAME.$ZFS_MODE. (If the property already exists, such a value will be added to the existing property if it is not already present.) If UPDATE is 'y', then, immediately upon successful backup, the current backup's snapshot will replace any existing snapshot that has the same property substring in one of two ways: either the existing snapshot will be removed (if NDMP service-generated) or its property will be cleared (if user-supplied). If the backup fails, the behavior will be the same as for UPDATE='n'.

If UPDATE is 'n', then any existing snapshot that contains $LEVEL.$DMP_NAME.$ZFS_MODE within its property will retain that property setting and (if an NDMP service-generated snapshot) will not be removed. Instead, the current backup's snapshot will be made inaccessible to future backups in one of two ways: either the snapshot will be removed (if NDMP service-generated) or its property will be cleared of the $LEVEL.$DMP_NAME.$ZFS_MODE value (if user-supplied).

The default value of UPDATE is 'y'. It is expected that this is the normative value as it allows snapshots for successive L0 through L8 backups to remain in the system to be used for a subsequent (L#+1) incremental backup. An administrator might want to set UPDATE to 'n' on an as-needed basis when he wants to perform a backup without upsetting the current set of snapshots.

**CAUTION**: Take care not to prematurely remove user-supplied snapshots (or any NDMP service-generated snapshots, for that matter) if the UPDATE variable is set and further incremental backups are planned. If you attempt an incremental backup that is dependent on a removed snapshot, a failure will result.

### Incremental Restore

To restore an incremental backup follow these guidelines:

- All prior level snapshots/backups will need to be applied (by the administrator) in addition to the specified level when a restore needs to be consistent with a given level.
- For L0, a nonexistent dataset must be specified.
- For L1-9, the dataset used for restoring L0 must be specified.
- If an incremental stream is received, then the destination file system must already exist, and its most recent snapshot must match the incremental stream's source.

If the administrator makes a modification to the dataset after restoring a given level, and then attempts to restore the next level from backup media, the attempt will normally fail because the dataset used for restore is no longer in the state required to receive the next level.

If this happens, two options are available to the administrator. Set the NDMP environment variable ZFS_FORCE or set the "Rollback datasets before restore" option in the NDMP Service configuration.

If this happens, two options are available to the administrator. Set the NDMP environment variable ZFS_FORCE or set the "Rollback datasets before restore" option in the NDMP Service configuration.

Backup Application Configuration

The following sections give guidance on how to configure various backup applications to use the NDMP service on the Sun ZFS Storage Appliance. Further details on application specifics can be found in the respective administration guides for the individual backup applications.

## Oracle Secure Backup

The following sections outlines the basic configuration steps needed to configure Oracle Secure Backup for NDMP operations.

Further information can be found in the *Oracle Secure Backup Adminstrator's Guide*. The configuration steps are shown using the Command Line utility **obtool**. Further information on **obtool** can be found in the *Oracle Secure Backup Reference Guide*.

### Sun ZFS Storage configuration

To ensure correct operations with Oracle Secure Backup and locally attached tape drives, the Sun ZFS Storage Appliance should have the NDMP service setting **DMA Tape Mode** set to **BSD**. This will ensure that the tape operations are correct for Oracle Secure Backup restore operations.

### Add NDMP host

To add a Sun ZFS Storage Appliance to perform normal dump backups

```
ob> mkhost -a ndmp -o -r mediaserver,client -i 10.9.176.23 -u admin -B dump -q aie7000
Password: ****
Password (again): ****
```

To add a Sun ZFS Storage appliance to perform ZFS-NDMP backups

```
ob> mkhost -a ndmp -o -r mediaserver,client -i 10.9.176.23 -u admin -B zfs -q aie7000
Password: ****
Password (again): ****
```

Verify the configuration

```
ob> lshost -l aie7000
aie7000-01:
    Access mode:            NDMP
    IP names:               10.9.176.23
    NDMP port:              (default)
    NDMP user name:         root
    NDMP password:          (set)
    NDMP backup type:       dump
    NDMP protocol version:  (default)
    NDMP auth type:         (default)
    Algorithm:              aes192
    Encryption policy:      allowed (Encryption supported only with encrypting drives)
    Rekey frequency:        1 month (system default)
    Key type:               transparent
    In service:             yes
    Roles:                  mediaserver,client
    Trusted host:           yes
    UUID:                   29b36f90-8eda-102d-a499-00144f9a7fb2
```

### Configure Storage Devices

Discover the locally-attached tape library and drives. With the **DMA Tape Mode** set to **BSD,** verify that the /dev/rmt device files returned are /dev/rmt/*bn.

```
ob> discoverdev -h aie7000 -v
Info: beginning device discovery for aie7000.
Info: connecting to aie7000

Info: devices found on aie7000:
   Info: SL500           ...
      Info: /dev/scsi/changer/c0t500104F000B966C3d0  attrs= [none]
         Info: WWN: [none]
         Info: SN:  559000203429
   Info: ULTRIUM-TD4     ...
      Info: /dev/rmt/0bn  attrs= norewind
         Info: WWN: [none]
         Info: SN:  1310204754
   Info: ULTRIUM-TD4     ...
      Info: /dev/rmt/1bn  attrs= norewind
         Info: WWN: [none]
         Info: SN:  1310218075

   aie7000-01_c0t500104F000B966C3d0  (new library)
      WWN: [none]
      new attach-point on aie7000-01, rawname /dev/scsi/changer/c0t500104F000B966C3d0

   aie7000-01_0bn  (new drive)
      WWN: [none]
      new attach-point on aie7000-01, rawname /dev/rmt/0bn

   aie7000-01_1bn  (new drive)
      WWN: [none]
      new attach-point on aie7000-01, rawname /dev/rmt/1bn
```

Associate the tape drive(s) with the tape library

```
ob> chd -t library -o aie7000_c0t500104F000B966C3d0
ob> chd -t tape -o -d 1 -l aie7000_c0t500104F000B966C3d0 aie7000-01_0bn
ob> chd -t tape -o -d 2 -l aie7000_c0t500104F000B966C3d0 aie7000-01_1bn

ob> lsdev
library             aie7000_c0t500104F000B966C3d0 in service
  drive 1           aie7000_0bn   in service
  drive 2           aie7000_1bn   in service
```

Inventory the library and label the tapes. The example labels the first 4 tapes

```
ob> inventory -L aie7000_c0t500104F000B966C3d0
ob> lsvol -L aie7000_c0t500104F000B966C3d0
Inventory of library aie7000_c0t500104F000B966C3d0:
    in    1:             barcode LT4016
    in    2:             barcode LT4005
    in    3:             barcode LT4006
    in    4:             barcode LT4007
    in    9:             barcode LT4015
    in    10:            barcode LT4008
    in    11:            barcode LT4009
    in    17:            barcode LT4017
```

```
    in    18:            barcode LT4018
    in    23:            barcode LT4019
    in    iee1:          barcode LT4002
    in    iee2:          barcode LT4001
    in    iee3:          barcode LT4000

ob> labelvol -D aie7000_0bn -f 1-4

ob> lsvol -L aie7000_c0t500104F000B966C3d0
Inventory of library aie7000-01_c0t500104F000B966C3d0:
    in    1:             barcode LT4016, 781243392 kb remaining, mediainfo hw
encryptable
    in    2:             barcode LT4005, 781243392 kb remaining, mediainfo hw
encryptable
    in    3:             barcode LT4006, 781243392 kb remaining, mediainfo hw
encryptable
    in    4:             barcode LT4007, 781243392 kb remaining, mediainfo hw
encryptable
    in    9:             barcode LT4015
    in    10:            barcode LT4008
    in    11:            barcode LT4009
    in    17:            barcode LT4017
    in    18:            barcode LT4018
    in    23:            barcode LT4019
    in    iee1:          barcode LT4002
    in    iee2:          barcode LT4001
    in    iee3:          barcode LT4000
```

### Create Backup Dataset

To create Backup Dataset for a dump backup

```
ob> mkds aie7000-backup1

include host aie7000

include path /export/data01
```

To create Backup Dataset for a zfs backup

```
ob> mkds aie7000-backup1

include host aie7000

include path /pool_0/local/hr/data01
```

### Perform a Backup

```
ob> backup -D aie7000-backup -l 0 --go
```

Use the **lsjob** and the **catxcr** commands in obtool to monitor the backup progress

### Perform a Restore

```
ob> restore -h aie7000 -p 1 -s latest --go /export/data01 -a /export/data01-restore
```

Use the **lsjob** and the **catxcr** commands in obtool to monitor the restore progress

## Symantec NetBackup

The following sections outlines the basic configuration steps needed to configure Symantec NetBackup for NDMP operations.

Further information can be found in the *Symantec NetBackup Adminstrator's Guide* as well as the *Symantec NetBackup for NDMP Adminstrator's Guide*.

### Add NDMP host

To configure the Sun ZFS Storage Appliance as a NDMP host

1. Open **NetBackup Administration Console**.
2. Expand **(Master Server)**, expand **Media and Device Management**, expand **Credentials**.
3. Right-Click **NDMP Hosts**, Click **New**.
4. Enter the Fully Qualified Domain Name (FQDN) or IP Address for the Sun ZFS Storage Appliance.

5. Click the radio button to enter the DMA Username and Password that was configured on the appliance.
6. Click OK.

### Configure Storage Devices

To configure locally attached tape devices

1. Open the **NetBackup Administration Console**. Select **Media and Device Management**.
2. In the right window pane, select **Configure Storage Devices**.
3. Click **Next** in the wizard.
4. Select the **Master Server**. Click **Change**. Enable **NDMP Server**, click **OK**.
5. Click **Next**.
6. Select the **NDMP Host** (Sun ZFS Storage Appliance) added above, click **Next**.
7. Click **Next** when wizard is completed with scan.
8. Click **Next** to bypass SAN Clients.
9. Review results, Click **Next**.
10. Select Devices that are intended for NetBackup use, click **Next**.

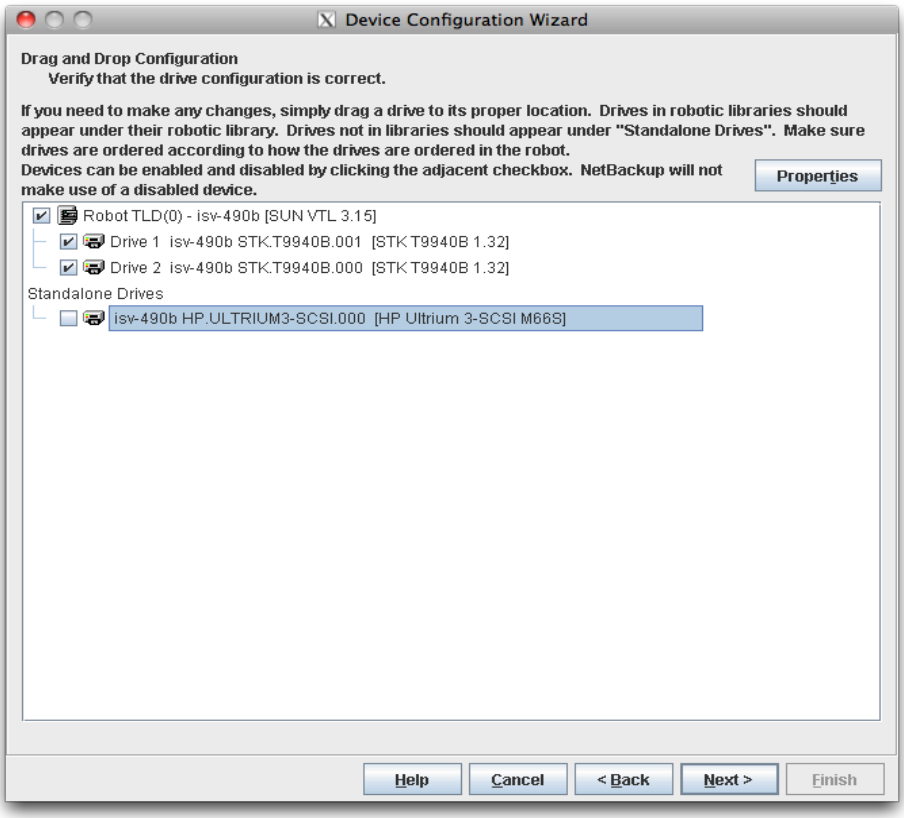
Device Configuration Wizard
Drag and Drop Configuration
Verify that the drive configuration is correct.
If you need to make any changes, simply drag a drive to its proper location. Drives in robotic libraries should appear under their robotic library. Drives not in libraries should appear under "Standalone Drives". Make sure drives are ordered according to how the drives are ordered in the robot.
Devices can be enabled and disabled by clicking the adjacent checkbox. NetBackup will not make use of a disabled device.
Properties
Robot TLD(0) - isv-490b [SUN VTL 3.15]
Drive 1 isv-490b STK.T9940B.001 [STK T9940B 1.32]
Drive 2 isv-490b STK.T9940B.000 [STK T9940B 1.32]
Standalone Drives
isv-490b HP.ULTRIUM3-SCSI.000 [HP Ultrium 3-SCSI M66S]
Help
Cancel
< Back
Next >
Finish

11. Follow wizard to add **Storage Unit** and complete wizard.

### Inventory Robot

1. Open **NetBackup Administration Console**. Expand **Media and Device Management**, expand **Devices**, right click **Robots**, select **Inventory Robot**.

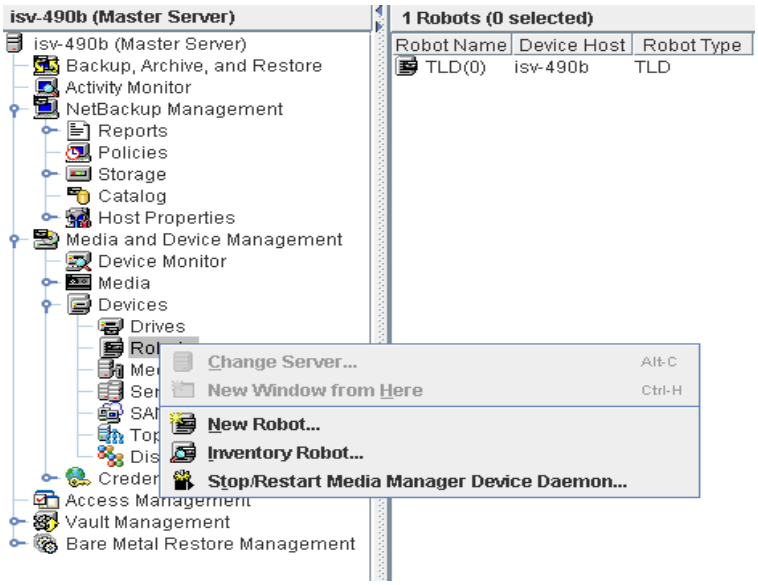


2. Select the **Device Host** and select the **Robot**.
3. Click the **Update Volume Configuration** radio button.
4. Click **Start**.
5. If prompted to **Update Volume Configuration**, select **Yes**.

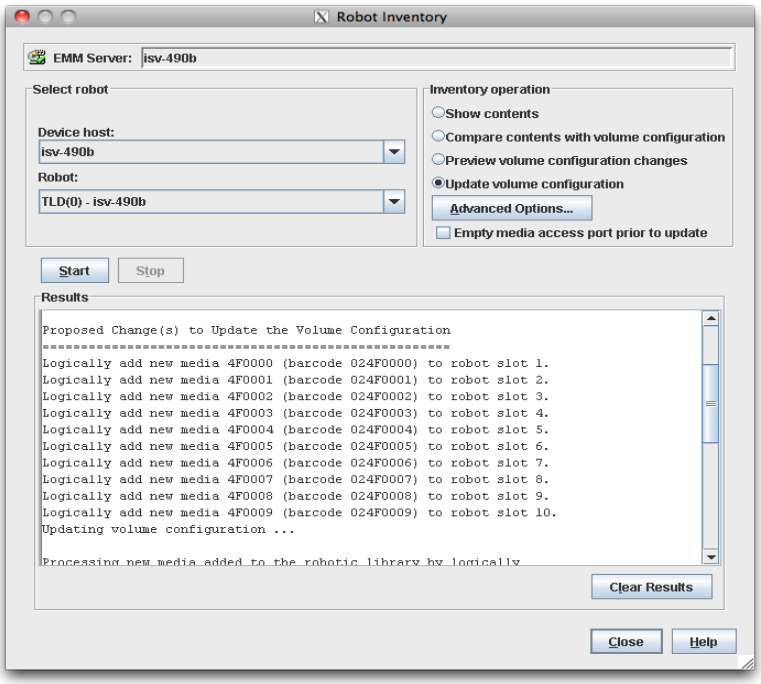
Robot Inventory
EMM Server: isv-490b
Select robot
Device host:
isv-490b
Robot:
TLD(0) - isv-490b
Inventory operation
Show contents
Compare contents with volume configuration
Preview volume configuration changes
Update volume configuration
Advanced Options...
Empty media access port prior to update
Start
Stop
Results
Proposed Change(s) to Update the Volume Configuration
Logically add new media 4F0000 (barcode 024F0000) to robot slot 1.
Logically add new media 4F0001 (barcode 024F0001) to robot slot 2.
Logically add new media 4F0002 (barcode 024F0002) to robot slot 3.
Logically add new media 4F0003 (barcode 024F0003) to robot slot 4.
Logically add new media 4F0004 (barcode 024F0004) to robot slot 5.
Logically add new media 4F0005 (barcode 024F0005) to robot slot 6.
Logically add new media 4F0006 (barcode 024F0006) to robot slot 7.
Logically add new media 4F0007 (barcode 024F0007) to robot slot 8.
Logically add new media 4F0008 (barcode 024F0008) to robot slot 9.
Logically add new media 4F0009 (barcode 024F0009) to robot slot 10.
Updating volume configuration ...
Clear Results
Close
Help

6. Click **Close**.

### Create Backup Policy

1. In the **NetBackup Administration Console**, expand **NetBackup Management**, right-click **Policy**, click **New Policy**.
2. Enter a name for new the policy.
3. Select the **Use Backup Policy Configuration Wizard** checkbox, click **OK**.
4. Select Policy Type **NDMP**, and then click **Next**.

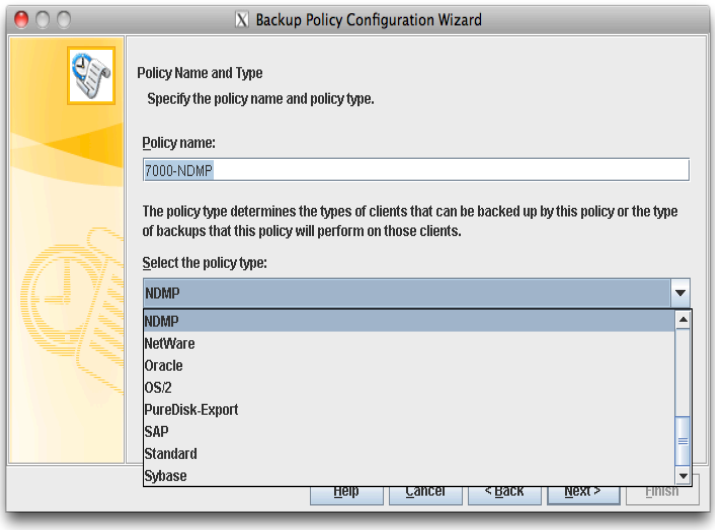


5. Click **Add** and enter hostname or IP address of Sun ZFS Storage appliance.
6. In the **Hardware and operating system** dialog, select **NDMP, NDMP**. Click **Add**.

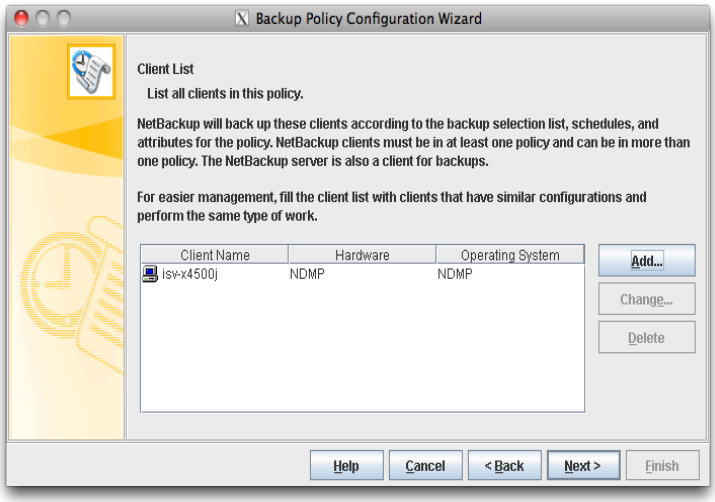
Backup Policy Configuration Wizard
Client List
List all clients in this policy.
NetBackup will back up these clients according to the backup selection list, schedules, and attributes for the policy. NetBackup clients must be in at least one policy and can be in more than one policy. The NetBackup server is also a client for backups.
For easier management, fill the client list with clients that have similar configurations and perform the same type of work.
Client Name
Hardware
Operating System
isv-x4500j
NDMP
NDMP
Add...
Change...
Delete
Help
Cancel
< Back
Next >
Finish

7. Click **Close**, Click **Next**.
8. Click **Add**, enter in backup path for share. Click **Next**.

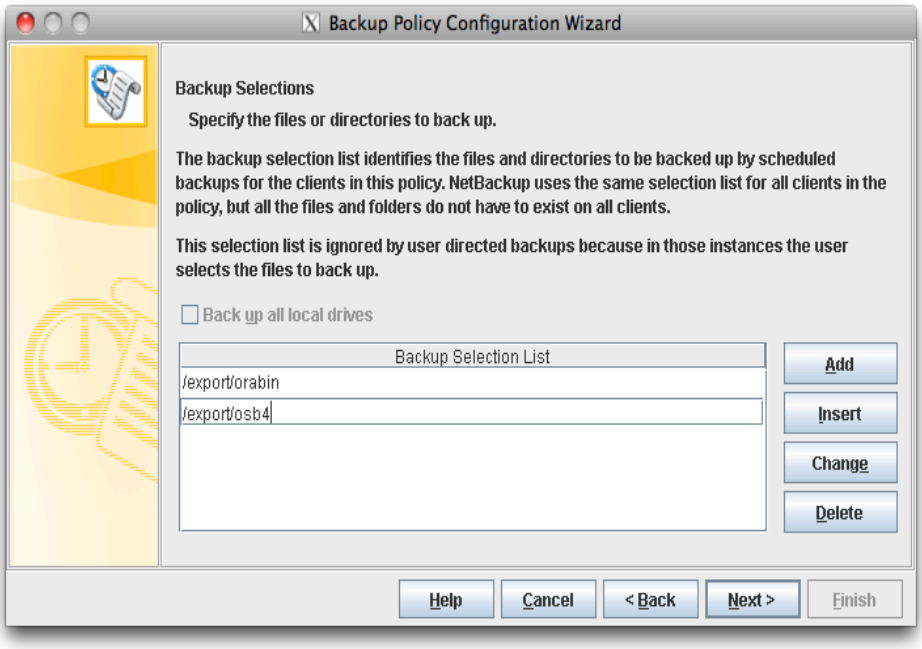
Backup Policy Configuration Wizard
Backup Selections
Specify the files or directories to back up.
The backup selection list identifies the files and directories to be backed up by scheduled backups for the clients in this policy. NetBackup uses the same selection list for all clients in the policy, but all the files and folders do not have to exist on all clients.
This selection list is ignored by user directed backups because in those instances the user selects the files to back up.
Back up all local drives
Backup Selection List
/export/orabin
/export/osb4
Add
Insert
Change
Delete
Help
Cancel
< Back
Next >
Finish

For a ZFS-NDMP backup, the Backup Selection List should include the variable **SET TYPE=zfs** as the first entry.

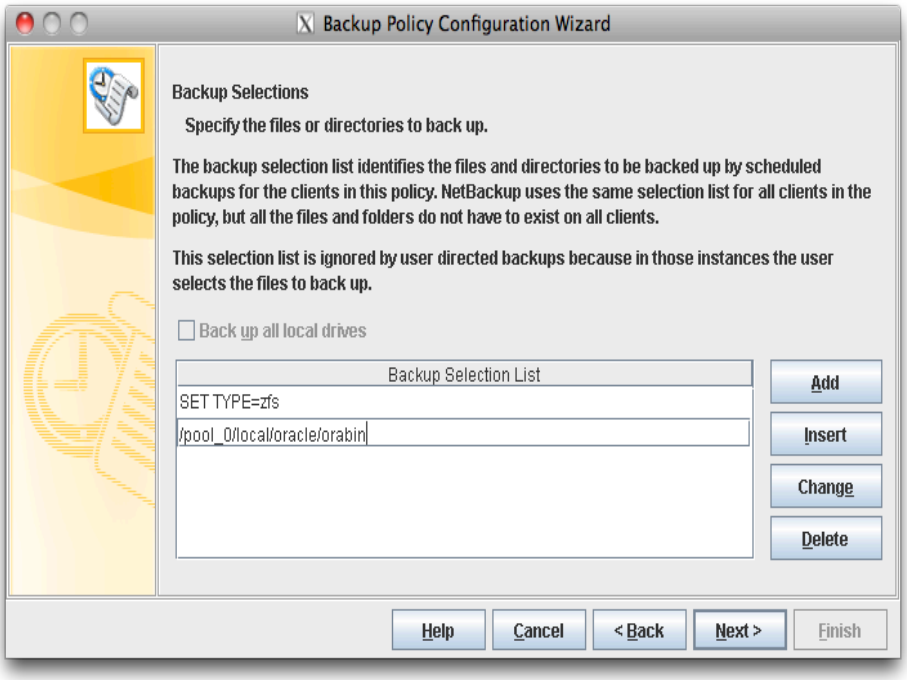
Backup Policy Configuration Wizard
Backup Selections
Specify the files or directories to back up.
The backup selection list identifies the files and directories to be backed up by scheduled backups for the clients in this policy. NetBackup uses the same selection list for all clients in the policy, but all the files and folders do not have to exist on all clients.
This selection list is ignored by user directed backups because in those instances the user selects the files to back up.
Back up all local drives
Backup Selection List
SET TYPE=zfs
/pool_0/local/oracle/orabin
Add
Insert
Change
Delete
Help
Cancel
< Back
Next >
Finish

9. Set Full and Incremental configuration and click **Next**.
10. Set Schedule and Retention,and click **Next**.

11. Set Backup window settings and click **Next**.
12. Click **Finish**.
13. From the **NetBackup Administration Console**, expand **NetBackup Management**, expand **Policy**, and then double-click on the policy you just created.
14. In **Policy Storage unit / lifecycle policy** drop-down menu, choose the storage unit that contains the devices attached to the appliance.
15. Select **Allow Multiple data streams** checkbox if the Backup Selection List contains multiple shares to back up.

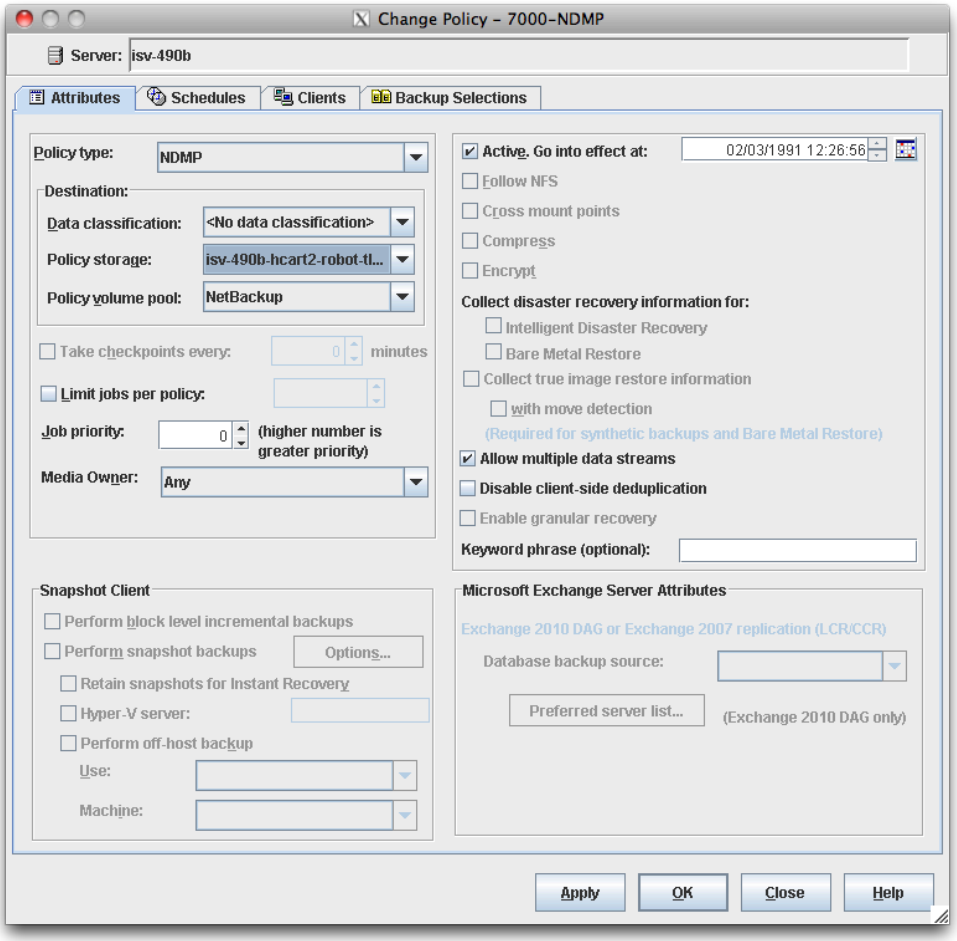


### Set Maximum Jobs per Client

If configured backup policies contain multiple shares and Allow multiple data streams is selected then the Global setting for Maximum jobs per client needs to be increased.

1. In **NetBackup Administration Console**, expand **NetBackup Management**, Expand **Host Properties**, click **Master Servers**.
2. Double-click the appropriate Master Server.
3. Under **Properties**, click **Global Attributes**.

4. In **Maximum jobs per client box**, increase the value to at least 2. To ensure all drives are used, the value should be equal to the number of tape drives connected to the appliance.
5. Click **OK** to finish.

## EMC Legato Networker

The following sections outlines the basic configuration steps needed to configure EMC Legato Networker for NDMP operations.

Further information can be found in the *EMC Legato Networker Administrator Guide*.

### Configuring the Sun ZFS Storage Appliance as a Storage Node

1. Run **jbconfig** from command line of Networker host.
2. Choose Option #3 Configure **an Autodetected NDMP SCSI Jukebox**.
3. Enter name of Sun ZFS Storage appliance that has connection to tape library and drives.
4. Enter NDMP username and password that has been configured in the NDMP service on the Sun ZFS Storage Appliance.
5. Follow the prompts to complete discovery.

After the library and drives are added via **jbconfig** Networker automatically performs an inventory on the library.

### Label Tapes

1. Log into **Networker Management Console**.
2. Choose appropriate **Networker Server**.
3. Select **Networker** from Managed Applications list.
4. Choose **Devices** from top panel.
5. Expand **Libraries**, choose appropriate library.
6. In far right-hand panel select media to label, right-click, choose appropriate **Target Media Pool**, uncheck **Prompt to Overwrite Existing Label**, select **Label**.

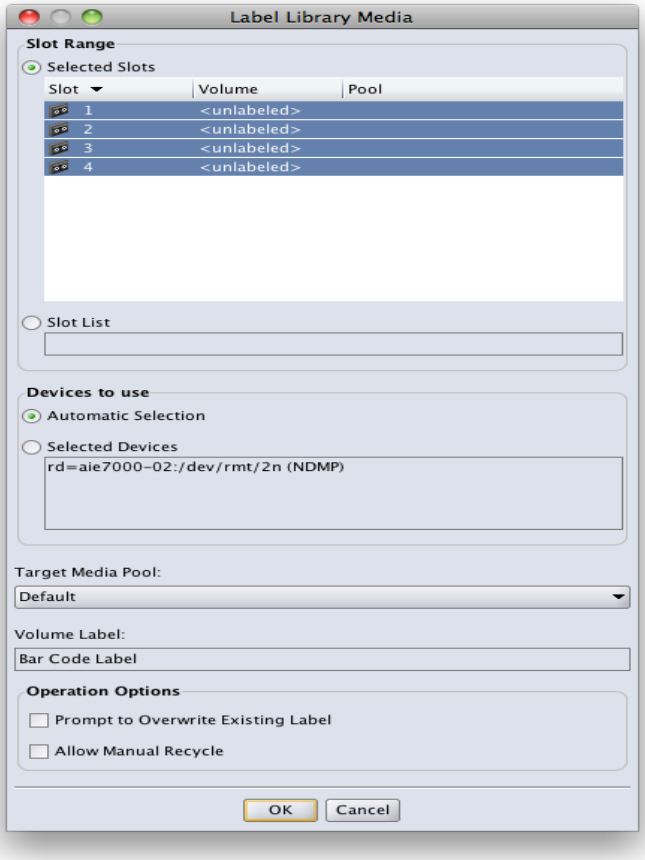


7. Click **OK**.

**CAUTION**: The **Prompt to Overwrite Existing Label** option should only be unchecked if the administrator is certain there is no foreign labeled media in the library that needs to be preserved.

### Add NDMP Host

1. Log into Networker Management Console.
2. Choose appropriate **Networker Server**.
3. Select **Networker** from Managed Applications list.
4. Choose Configuration from top panel.
5. Choose Clients.
6. Click Green + symbol to add new client.
7. Select **General Tab**. In **Identity** section enter name or IP Address of Sun ZFS Storage appliance. In **Save set** section enter the appropriate backup path(s).

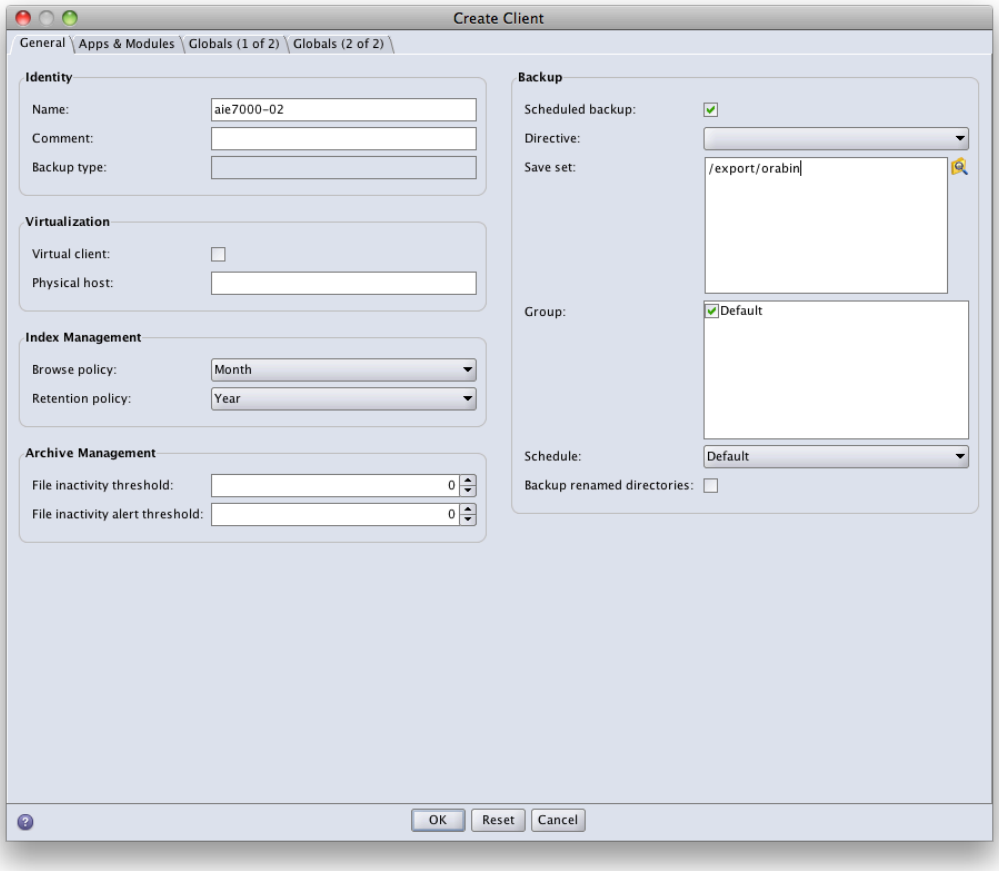


8. Select **Apps & Modules Tab**. In the **Access** section enter the Remote User and **Password** that have been configured in the NDMP service for the Sun ZFS Storage appliance.
9. In the Backup section, enter the Backup command **nsrndmp_save -T dump**, then check the **NDMP** option box. Enter **NDMP array name** as the name or IP address of the Sun ZFS Storage appliance.
10. In the Application Section. enter **HIST=y, UPDATE=y, DIRECT=y**.

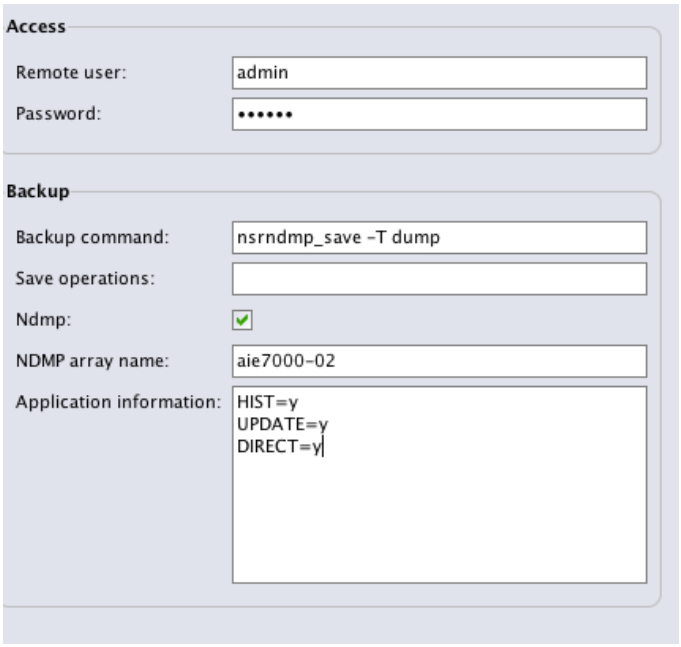


11. Select the **Globals (2 of 2) tab**. Enter the name of the Sun ZFS Storage appliance as **Storage node**.

### Setup for ZFS-NDMP backups

To configure EMC Legato Networker for ZFS-NDMP backups enter the Save Set information in the form of a valid ZFS dataset path.

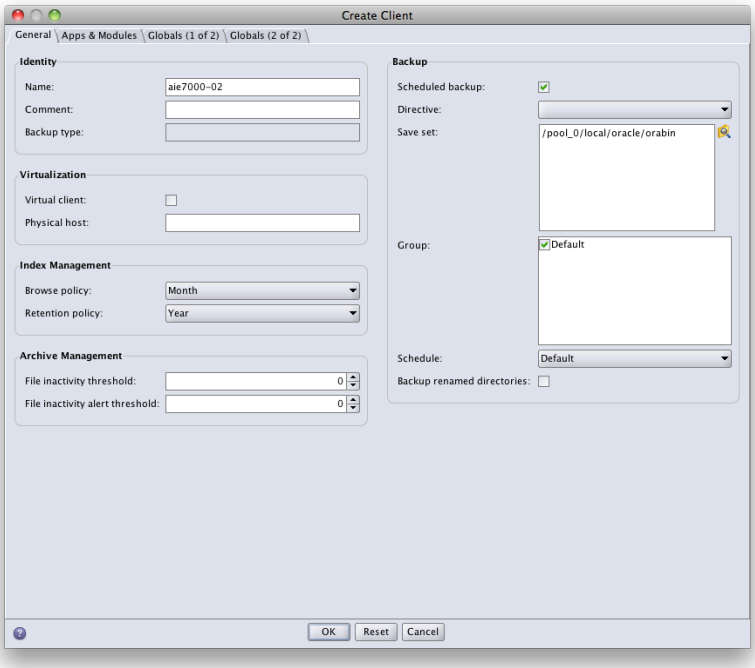

Additionally, make the following changes on the **Apps and Modules** tab.

1. **The Backup command** should be **nsrndmp_save -T zfs.**
2. The Application information section should be left empty.

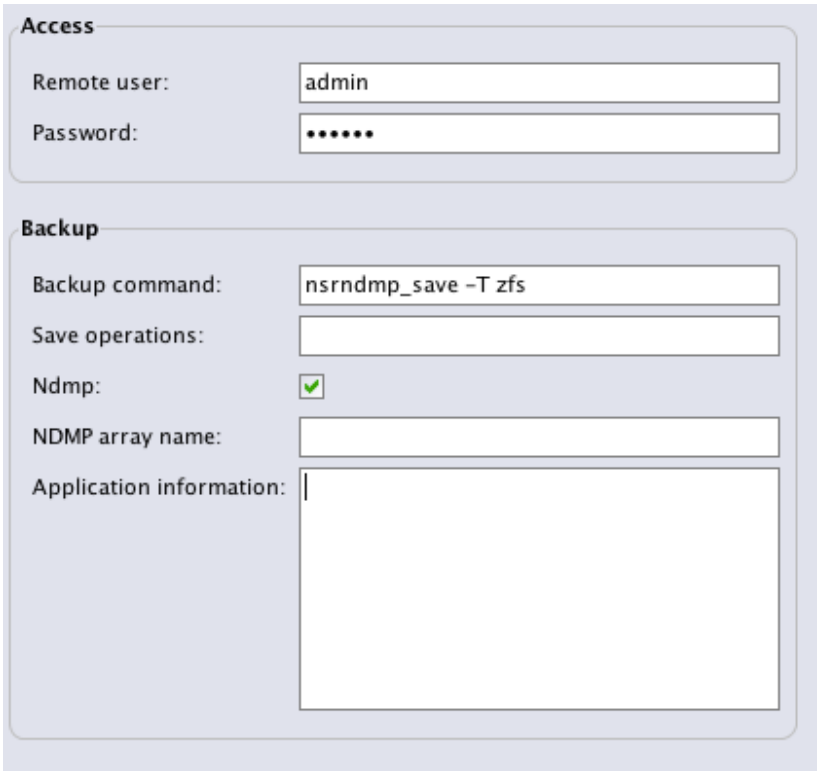

## Conclusion

The Sun ZFS Storage Appliance supports the industry standard Network Data Management Protocol (NDMP) to ensure efficient data protection for backup and recovery. The Sun ZFS Storage appliance works with a variety of industry-leading backup applications (such as Oracle Secure Backup, Symantec NetBackup, and EMC Legato Networker) to offer a complete solution for data protection needs.

NDMP Implementation Guide
for the Sun ZFS Storage Appliance
December 2011
Author: Ryan Arneson

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

**Hardware and Software, Engineered to Work Together**